A Multiple Hypothesis Tracker for a Distributed Network of Sensors

Constantino Rago Charlie Gershfield
BAE Systems
Advanced Information Technologies
Fusion Technology and Systems Division
Burlington, MA, USA

Thomas M. Strat
Defense Advanced Research Projects Agency (DARPA)
Information Exploitation Office

Abstract - ***This paoer describes the Multiple Hypothesis Network Tracker (MHNT) being developed by BAE Systems. The goal here was to design, develop and implement a tracker capable of tracking multiple vehicles moving across a network of geographically distributed static sensors, where each sensor's coverage area is small compared to the distance between sensors. Each sensor reports the time when a target crosses the sensor's coverage area Sensors may also report features associated with the target. A centralized tracker receives reports from all of the sensors and, using prior information regarding target motion across the network of sensors and actual travel times, assign reports to tracks and computes the most likely set of assignments.***

Keywords: Tracking, Network of Sensors, Distributed Sensors, Multiple Hypothesis Tracking.

1 Introduction

In the problem at hand, we are interested in tracking multiple targets as they move across a given area of interest (AOI). Targets move along roads, and the locations of the sensors with respect to the road are known. Sensors are sparsely distributed across the AOI. They are fixed and static -they can not move or look into a particular area. The sensor's coverage area always stays the same, and they can only detect a target whenever the target crosses their coverage area. They are also able to detect the travelling direction of the target. Each sensor generates a report whenever a target leaves its coverage area, and these reports are sent to a central processor running the MHNT. A report consists of a time stamp, a sensor ID (that also encodes the target travelling direction, as discussed later) and features extracted from the target.

The network tracking problem can be seen as the dual of the standard report-to-track tracking problem encountered in radar, sonar, video, and other sensor domains. In the standard tracking problem, the sensor observes the location of a target at a given time, while in the network tracking problem, the sensor observes the time a target crosses a given location. The fact that time is the observed variable leads to a formulation where the target state is time instead of position. To further emphasize this duality, in the standard tracking problem there is uncertainty in the location of the target at a given time, while in our network tracking problem formulation there is uncertainty in the time the target will cross a given location.

The tracking problem in this case is mainly an association problem: deciding which reports go with which targets. Dynamic target models are of little use due to the large time period between reports from the same target with respect to the target dynamics; for example, a target can stop and restart multiple times between observations. The target state consists of the time and location it was last seen, the direction it was travelling, and a feature state associated with the target. Given the fixed nature of the sensors, as well as the prior knowledge about roads and sensor locations, we can represent all this information using a directed graph (see section 2). The time it takes a target to travel from one sensor to another (which we will denote as the transit time) can be modelled as a random variable with a known prior distribution. Multiple hypothesis regarding report to track associations are kept by the tracker, and this prior information is then used, in conjunction with the actual travel times and measured features, to decide on the best set of report-to-track associations.

The paper is organized as follows: section 2 describes how to model prior information using a connectivity graph. Section 3 combines prior information with measurements to derive the posterior probability of a global hypothesis. Section 4 extends the global hypothesis posterior probability computation to include false alarms and missed detections. Section 5 describes the technique used to maintain multiple global hypotheses in real time. The last section (6) deals with the use of a scoring mechanism between reports.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **JUL 2006**	2. REPORT TYPE	3. DATES COVERED **00-00-2006 to 00-00-2006**

4. TITLE AND SUBTITLE **A Multiple Hypothesis Tracker for a Distributed Network of Sensors**	5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Defense Advanced Research Projects Agency (DARPA),Information Exploitation O?ce,3701 North Fairfax Drive,Arlington,VA,22203**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
9th International Conference on Information Fusion, 10-13 July 2006, Florence, Italy. Sponsored by the International Society of Information Fusion (ISIF), Aerospace & Electronic Systems Society (AES), IEEE, ONR, ONR Global, Selex - Sistemi Integrati, Finmeccanica, BAE Systems, TNO, AFOSR’s European Office of Aerospace Research and Development, and the NATO Undersea Research Centre.

14. ABSTRACT
see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**8**	

2 The observer connectivity graph

In the MHNT formulation, all prior information regarding roads, sensor locations, target motion, traffic patterns, and sensor performance is represented through the observer connectivity graph. The observer connectivity graph is a directed graph consisting of nodes and links that connect two nodes. Being a directed graph, each link has a direction associated with it.

Sensors are represented by one or more nodes in the graph. A sensor and a travelling direction in the sensor's coverage area are represented by a unique node. This node is called an "observer". As a target moves along the roads, from one sensor location to another, following a particular road in a given direction, in our mathematical model the target moves from one observer to another "through" a particular directed link. For the tracker, reports are generated by observers, not sensors, therefore there is an implicit location and travelling directions associated with each report. The tracker itself does not know anything about sensors, roads and traffic; it only knows about observers, links, and parameters or functions associated with those. In our current implementation, all parameters associated with the observer graph are estimated online from measured traffic densities and estimated transit times (derived from report associations produced by the tracker).

2.1 Observer characterization

As mentioned before, each observer is associated with a travelling direction in a sensor's coverage area. In order to characterize the sensor's performance in terms of detecting targets crossing its coverage area, two parameters are needed: the false alarm rate and the probability of missed detection. The false alarm rate, in targets per second, is the number of reports per unit time generated by the observer in the absence of real targets. The probability of missed detection is the probability that the sensor will miss reporting on a real target crossing its coverage area.

The other parameter associated with an observer has to do with traffic characterization: the observer's birth rate. The birth rate is the number of new targets per second expected to be seen by the observer. New targets are defined as those targets that have not been seen by any other observer previously. The observer's birth rate parameter provides a means of estimating the rate at which targets enter the AOI, e.g.: entry gates, roads that are not being monitored, parking lots, etc.

2.2 Link characterization

In the observer connectivity graph, nodes are connected to each other via directional links A link between two observers in the graph represents a piece of road (and a travelling direction on that road) between two sensors. A link between two observers represents the fact that a target can travel from the up-stream observer to the down-stream observer without crossing any other observer's coverage area. Generally, a link may represent more than one actual path between the observers. A target travelling through this link takes a certain amount of time to go from one end to the other (up-stream observer to down-stream observer). This travel time will be represented by a random variable with a given probability density function (pdf). This pdf represents the behavior of all targets instead of being associated with a particular target. The other parameter associated with a link is its probability. Given a non-trivial road network, there are, in general, multiple roads a target can take after crossing a particular sensor. The link's probability represent the fraction of the entering traffic that goes through the link and cross the sensor's coverage area at the other end of the link. The probability that a target will stop after having gone through an observer and before crossing any other observers is equal to:

$$P\{\text{stop after obs. } i\} \quad = \quad \bar{p}_i = 1 - \sum_{\forall j} P\{\text{link}_{i,j}\} \quad (1)$$

where the sum is over all links originating at the i^{th} observer. The stop probability represents the probability that a target may stop along the road, may enter a parking lot, may exit the AOI, etc.

An example of a network of sensors, roads and the corresponding observer connectivity graph is shown in Figure 1. In this case, there are three sensors, Sensor 1, Sensor 2 and Sensor 3 , and for each of these we define two observers, one for each travelling direction. The associated observer connectivity graph has 6 nodes, one for each observer, as shown in the bottom part of the same figure.

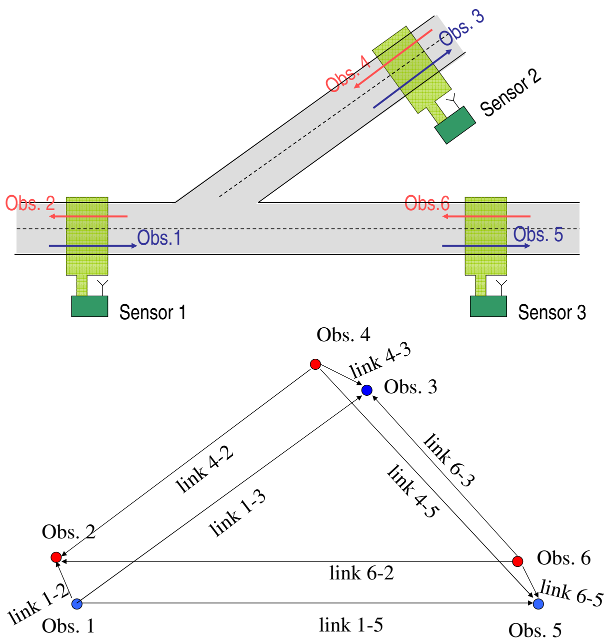

Figure 1: A simple example illustrating the relationship between sensors, roads, and the observer connectivity graph.

3 Global hypothesis posterior probability

As mentioned before, the goal of the tracker is to assign reports to tracks in some optimal way. In the case of a multiple hypothesis tracker, we will evaluate and keep multiple feasible assignments. A set of assignments such that every report gets assigned to one and only one track is usually denoted as a global hypothesis. Due to the combinatorial nature of the problem, and in order to have a real-time tracker, there is a need to limit the number of feasible assignments for a given report. The assignments corresponding to the global hypothesis with the maximum posterior probability represents the best track estimates at the current time.

Conditioned on our observation set, we wish to choose the most probable global hypothesis. The posterior probability of a global hypothesis will be denoted by $P\{H_i|Y\}$, where H_i is a feasible report-to-track assignment global hypothesis, and Y is the set of all reports received up to a particular time ($Y : \{R1, R2, ..., RN\}$). Although not explicitly shown, prior information regarding traffic, roads, and camera locations -as contained in the observer graph- is assumed to be known and available.

We start by re-writing the posterior probability of a global hypothesis in terms of the *a-priori* probabilities of all feasible global hypotheses $P\{H_j\}$. The posterior probability of the i^{th} global hypothesis, H_i, given the observation set Y is:

$$P\{H_i|Y\} = \frac{P\{H_i, Y\}}{P\{Y\}} = \frac{P\{Y|H_i\}P\{H_i\}}{P\{Y\}} \tag{2}$$

Note that, in the absence of false alarms and missed detections, and due to the fact that there is no observation error, $P\{Y|H_i\} = 1$ for each H_i in the set of feasible global hypotheses (*i.e., a global hypothesis is a collection of individual tracks, and given a track, there is only one set of feasible measurements). Handling of false and missed detection will be included in section (4). The probability of getting a measurement set Y is:*

$$P\{Y\} = \sum_{all\ feasible\ H_j} P\{Y|H_j\}P\{H_j\} \tag{3}$$

Replacing eqn. 3 in eqn. 2, we get:

$$P\{H_i|Y\} = \frac{P\{H_i\}}{\sum_{all\ feasible\ H_j} P\{H_j\}} \tag{4}$$

The function of the tracker is to choose a value for i that maximizes the above expression. To do this, we must specify first how to compute the a-priori *probability of a global hypothesis. Next, we look at the individual events that form a global hypothesis, including the target birth probability (see subsection 3.1), the transit time between observers (subsection 3.2) and the probability of not having seen the target (subsection 3.3). We can then compute the prior probability (subsection 3.4) and the posterior probability (subsection 3.5) of a global hypothesis.*

3.1 Target birth probabilities

Let us assume that, for any given observer, the probability of k new target births in a time interval T is given by a Poisson distribution (see [1]) with parameter equal to the birth rate λ_B of that particular observer:

$$P\{k\ new\ targets\ in\ interval\ T\} = \frac{(\lambda_B\ T)^k}{k!}\ e^{-\lambda_B T} \tag{5}$$

The probability of a single target birth in an infinitesimally small time interval dt is given by:

$$P\{one\ new\ target\ in\ interval\ dt\} = (\lambda_B\ dt)\ e^{-\lambda_B\ dt} \tag{6}$$

Setting $k = 0$, the probability of no target births during a time interval $t_2 - t_1$ is given by:

$$P\{no\ new\ targets\ in\ interval\ (t_2 - t_1)\} = e^{-\lambda_B(t_2-t_1)} \tag{7}$$

We can now calculate the probability of the target births assumed by any global hypothesis by considering each observer separately. We let $\lambda_B(j)$ denote the target birth rate at the j^{th} observer. The global hypothesis assumes that certain reports are from new targets (target births). In particular, for j^{th} observer, the first birth occurs at time $t_{j,1}$, second birth at time $t_{j,2}$, and so on.

Given that our observers are reporting from time 0 to time t, and noting that target births at each observer are independent from one observer to the next, we can write the probability of the target birth scenario assumed by the i^{th} global hypothesis for all observers as follows:

$$\begin{aligned} &P\{birth\ scenario\ in\ i^{th}\ global\ hypothesis\} \\ &= \prod_{j=1}^{V} \Bigg[\underbrace{e^{-\lambda_B(j)t_{j,1}}}_{Time\ to\ first\ birth}\ \underbrace{\lambda_B(j)\ dt\ e^{-\lambda_B(j)dt}}_{first\ birth} \\ &\quad \prod_{k=2}^{N_{Bj}} \left(\underbrace{e^{-\lambda_B(j)(t_{j,k}-t_{j,k-1}-dt)}}_{time\ between\ births} \underbrace{\lambda_B(j)\ dt\ e^{-\lambda_B(j)dt}}_{k^{th}\ birth} \right) \\ &\quad \underbrace{e^{-\lambda_B(j)(t-t_{j,N_{Bj}}-dt)}}_{time\ from\ last\ birth} \Bigg] \end{aligned} \tag{8}$$

where

- *V : the number of observers,*
- *N_{Bj} : the number of target births assumed to have occurred at the j^{th} observer, occurring at times $t_{j,k}$, $k = 1 : N_{Bj}$*

The term $e^{-\lambda_B(j)dt}$ cancels out with the term $e^{+\lambda_B(j)dt}$, and given that the birth terms do not depend on k, we can rewrite the above as:

$$\begin{aligned} &P\{birth\ scenario\ in\ i^{th}\ global\ hypothesis\} \\ &\qquad = (dt)^{N_B} \prod_{j=1}^{V} \lambda_B(j)^{N_{Bj}}\ e^{-\lambda_B(j)\ t} \end{aligned} \tag{9}$$

where $N_B = \sum_{j=1}^{V} N_{Bj}$ denotes the total number of target births assumed by the i^{th} global hypothesis.

The next step is to model the probabilities associated with transit times between observers and the probability of not having seen the target since its last observation.

3.2 Transit time probabilities

We start by defining the pdf of the time at which the target passes observer j as a function of the time it passed observer i, as well as other road/traffic parameters. In particular, we assume that this density function is stationary -i.e., it does not depend on the specific times but instead depends on the time difference (transit time). We characterize the transit time density function between a pair (i, j) of observers by a shifter one-sided pdf, with a minimum transit time denoted by τ_{ij}^{min} and an average transit time denoted by τ_{ij}^{avr}.

The minimum transit time has a direct impact on the "early-gate" function, limiting the number of hypotheses generated by discarding "early" reports, while the average transit time has a direct impact on the value of the probability, as well as on the "late" gate. In the tracker implementation, we used the same pdf function (with different parameters) for all links. The function was selected based on a good fit to the transit time data we had collected, as well as the fact that it should have an easy to compute cumulative function (see subsection 3.3).

The pdf functions that models the transit time between observers i and j will be denoted by $f_{ij}(\tau)$, and the corresponding cumulative density function by $F_{ij}(\tau)$

The probability of a target travelling from observer j to observer i in (τ_k) is approximately given by:

$$P\{traveling\ from\ observer\ j\ to\ i\ in\ \tau_k\ sec\} \approx f_{ij}(\tau_k)dt \tag{10}$$

This equation is used to compute the total probability associated with the transit times postulated by a global hypothesis. Given a global hypothesis, let N_T be the number of tracks in the hypothesis. Suppose there are $R_j + 1$ reports assigned to the j^{th} track including the original report in which the target was born. Then, there are R_j transit times for the j^{th} track, and consequently, a total of $N_A = \sum_{j=1}^{N_T} R_j$ transit times in the global hypothesis. Assuming that the density function of each transit time is independent of the others, the probability of the observed transit times in the global hypothesis is:

$$\begin{aligned} P\{\ (\tau_1, \ldots, \tau_{N_A})\ transit\ times\} &= \prod_{n=1}^{N_A} f_{i_{n-1},i_n}(\tau_n)dt \\ &= (dt)^{N_A} \prod_{n=1}^{N_A} f_{i_{n-1},i_n}(\tau_n) \end{aligned} \tag{11}$$

where

- i_{n-1} and i_n denote respectively the observers of the reports $n-1$ and n,
- τ_n is the time difference between the times of the reports $n-1$ and n,
- $f_{i_{n-1},i_n}(\cdot)$ is the transit time pdf for the n^{th} report assignment with appropriate parameters chosen based on the time of the $n-1^{th}$ report, the time of the n^{th} report and the link parameters between the observers i_{n-1} and i_n.

3.3 Probability of Not Seen Yet (NSY)

We now consider the probability of not having seen a particular target since the last observation, as assumed by a given global hypothesis. In particular, let observer i be the last observer that has seen the target according to a global hypothesis. Let observer j be located downstream from observer i in the observer connectivity graph. The probability of the target not seen yet (NSY) by observer j at the current time t is then given by

$$P\{target\ Not\ Seen\ Yet\ at\ time\ t\} = 1 - F_{ij}(t)$$

where $F_{ij}(\dot{})$ represents the cumulative density function of the transit time distribution between observers i and j. This probability is represented by the shaded area in Figure 2.

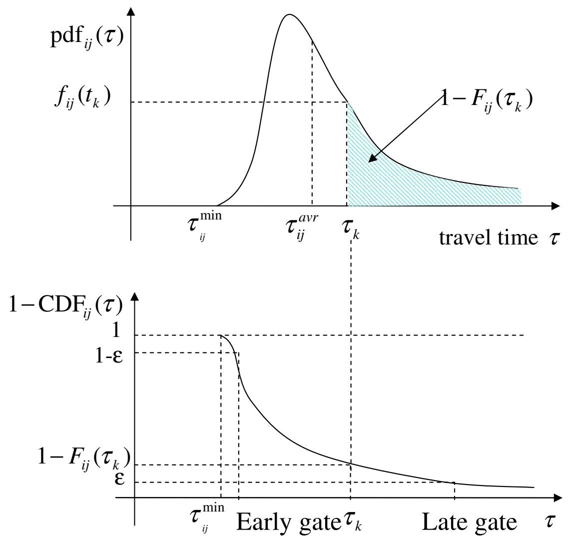

Figure 2: Top: Probability density function for transit time between observers i and j. Bottom: Probability of target "not seen yet" at time t_k by observer j.

By considering all the links the target could have followed after passing observer i, and the a-priori probability of the target following each of these links, as well as the probability of the target stoping/leaving before reaching any of the down-stream observers, we can write the probability of a target not having been seen yet at time t as follows:

$$P\{NSY(t)\} = \bar{p}_i + \sum_{j \in O_i} p_{ij}(1 - F_{ij}(t)) \tag{12}$$

where:

- $\bar{p}_i$ is the probability that the vehicle stopped or exited from the camera network after passing through observer i
- p_{ij} is the a-priori probability of the target following link (i, j)

- O_i is the set of all observers located 'down-the-road' (downstream) from observer i as defined by the links in the observer connectivity graph.

We note here that: $\bar{p}_i + \sum_{j \in O_i} p_{ij} = 1$, reflecting the fact that we have enumerated all the possible outcomes of a vehicle having passed the i^{th} observer.

We can now write the probability of not having seen any of the targets since their last observation by the system. Assuming that the global hypothesis contains N_T tracks and the current time is, this probability is given by:

$$
\begin{aligned}
&P\{\text{targets in a global hypothesis NSY at time } t\} \\
&\quad = \prod_{k=1}^{N_T} \left[\bar{p}_{l_k} + \sum_{j \in O_{l_k}} p_{l_k,j}(1 - F_{l_k,j}(t)) \right] \qquad (13)
\end{aligned}
$$

where

- l_k denote the observer that generated the latest report associated with track k,
- $F_{lk,j}(\cdot)$ represent the cumulative density function of the transit time from observer l_k to a downstream observer j.

3.4 Global hypothesis *a-priori* probability

We can now write the a-priori probability of the i^{th} global hypothesis by decomposing the global into the three components we have just covered: births, transit times and NSY's. Combining the probability of the birth scenario [equation (9)], the probabilities of the transit times corresponding to the hypotheses's report assignments [equation (11)], and the probability of not having seen the targets since their last sighting [equation (13)], we have:

$$
\begin{aligned}
P\{H_i\} \quad = \quad & \underbrace{(dt)^{N_B} \prod_{j=1}^{V} \lambda_B(j)^{N_{Bj}} e^{-\lambda_B(j)\, t}}_{\text{target births}} \\
& \underbrace{(dt)^{N_A} \prod_{n=1}^{N_A} f_{i_{n-1},i_n}(\tau_n)}_{\text{report assignments}} \\
& \underbrace{\prod_{k=1}^{N_T} \left(\bar{p}_{l_k} + \sum_{j \in O_{l_k}} p_{l_k,j}(1 - F_{l_k,j}(t)) \right)}_{\text{not seen yet}} \qquad (14)
\end{aligned}
$$

We note that the total number of reports, $R = N_A + N_B$, including those representing target births and additional report assignments, is the same for all global hypotheses. We can regroup the different terms in the above expression into a part that is the same for all global hypotheses and a part that is a function of the particular global hypothesis:

$$
\begin{aligned}
P\{H_i\} \quad = \quad & \underbrace{(dt)^{R} \prod_{j=1}^{V} e^{-\lambda_B(j)\, t}}_{\text{Hypothesis independent}} \\
& \prod_{j=1}^{V} \lambda_B(j)^{N_{Bj}} \prod_{n=1}^{N_A} f_{i_{n-1},i_n}(\tau_n) \\
& \prod_{k=1}^{N_T} \left(\bar{p}_{l_k} + \sum_{j \in O_{l_k}} p_{l_k,j}(1 - F_{l_k,j}(t)) \right) \qquad (15)
\end{aligned}
$$

3.5 Global hypothesis posterior probability

Recalling equation 4, we have:

$$
P\{H_i|Y\} = \frac{P\{H_i\}}{\sum_{j=1}^{N_G} P\{H_j\}} \qquad (16)
$$

where N_G is the total number of feasible global hypotheses. Letting

$$
C = (dt)^R \prod_{j=1}^{V} e^{-\lambda_B(j)(t)} \qquad (17)
$$

and

$$
\begin{aligned}
D_i \quad = \quad & \prod_{n=1}^{N_B} \lambda_B obs(n) \prod_{k=1}^{N_A} f_{i_{n-1},i_n}(\tau_n) \\
& \prod_{k=1}^{N_T} \left(\bar{p}_{l_k} + \sum_{j \in O_{l_k}} p_{l_k,j}(1 - F_{l_k,j}(t)) \right) \qquad (18)
\end{aligned}
$$

then posterior probability of the i^{th} global hypothesis conditioned on the measurements is given by:

$$
P\{H_i|Y\} = \frac{C\ D_i}{\sum_{m=1}^{N_G} C\ D_m} = \frac{D_i}{\sum_{m=1}^{N_G} D_m} \qquad (19)
$$

Dropping the denominator which is common to all feasible global hypotheses, the best global hypothesis H_i is selected based on the following condition

$$
D_i \geq D_j \qquad \forall j \qquad (20)
$$

4 False alarms and missed detections

In the first part of this section we extend the global hypothesis posterior probability calculations to include the possibility that a report is a false alarm, and therefore does not associated with any target. In order to include this possible explanation for a report, we need to first formulate a model for false alarm generation (observer dependent) and then derive the corresponding probabilities.

In the second part of this section we extend the global hypothesis probability calculations to include the fact that some of the observers may miss reporting a vehicle when it crosses their coverage area. In order to handle possible missed reports, the tracker should look at "incomplete" track hypothesis (incomplete in the sense that a track may 'skip' an observer). The probability of a global hypothesis that includes missed reports should be modified accordingly, as shown in subsection (4.2).

4.1 False alarms

Individual false alarm reports are not different from reports generated by a real target (if they were different, then it would be possible to filter them out). Is only when looking at the track level (a collection of reports) that the differences manifest: false alarms do not 'travel' across the network of sensors as a real targets do, so we expect false alarms to be uncorrelated in time.

Assuming that the false alarms generated by each observer can be modelled as a Poisson process uncorrelated across observers, we can the use target's birth model derived in section 3.1 to model the false alarm generation process. Denoting by $\{Y_1 : Y_{N_{FA}}\}$ the set of reports assumed to be false alarms by hypothesis H_i, then the probability of this set of false alarms is given by:

$$P\{\textit{false alarms}\} = (dt)^{N_{FA}} \prod_{j=1}^{V} \lambda_{FA}(j)^{N_{FA}(j)}\, e^{-\lambda_{FA}(j)\, t} \tag{21}$$

where

- N_{FA} : *number of false alarms postulated by hypothesis* H_i
- $\lambda_{FA}(j$ *:false alarm rate of observer j (in reports per second)*
- $N_{FA}(j)$: *number of false alarms assigned to observer j*

4.2 Missed detections

As mentioned above, sometimes an observer will fail to report a vehicle crossing its coverage area; for example, the vehicle is obscured by another vehicle, or two consecutive vehicles are reported as a single one due to the finite resolution of the sensor. Irrespective of the cause that makes the sensor fail to report, this phenomenon can be characterized by a probability of miss detection associated with each observer. Missed detections have two main implications in the formulation of the tracker:

- *when gating a report with existing tracks, the tracker should now look beyond the immediate up-stream observers. Until now, the tracker would gate a report against those tracks whose last report originated at the parent observers. Now , the tracker should look also among those tracks where the last report originated at the grandparents, grand-grandparents, and so on (assuming that a track could have multiple consecutive and nonconsecutive misses).*
- *when computing the probability that a target has not been seen yet by the "down-stream" observers, the tracker should consider the possibility that the vehicle has passed the down-stream observer, but the observer has missed to report it and the vehicle has not yet reached the next observer. Same logic should be extended to possible multiple down-stream misses.*

4.3 Global hypothesis posterior probability including missed detections

As mentioned before, the global hypothesis probability can be computed by evaluating the probability of each of the tracks included in that global. If the current report, reported by observer k at time t_k gates with a track whose last report corresponds to observer i at time t_i, and if there is no direct link between observers k and i, but there is a path involving multiple links and observers, then the contribution of the current report to the track probability is given by:

$$P\{\textit{report } t_k \textit{ continuation of track from report at } t_i\}$$
$$= f_{i,k}^{equiv}(t_k - t_i)\, p_{k,k-1} \prod_{j=k-1}^{i} p_{miss}(obs(j))\, p_{j,j-1} \tag{22}$$

It should be noted that there may be multiple paths between any two non-consecutive observers. If the report gates with a track through more than one path, then the path with the highest combined probability (eqn. 22) will be selected for that hypothesis.

4.3.1 Probability of NSY calculation in the presence of miss detections

Both gating and transit time probability calculations involve looking backwards (or up-stream) on the observer connectivity graph. To compute the probability of NSY, the tracker needs to look forward (or down-stream) on the graph. Equation (12) shows the probability of not having seen a track at time t without considering missed detections. When an observer may fail to report the target, then, in addition to considering that the target stopped after the last report's observer, or have not yet reached any of the children observers, we need to consider the possibility that the target did reach one of the children observers but was not reported, and either stopped after that or has not reached any of the grand children observers. In this case, the probability of not having-seen-yet the target should be modified as follows:

$$\begin{aligned} P\{NSY(t)\} &= \\ &\bar{p}_{obs(k)} + \sum_{j \textit{ (childs of } obs(k))} p_{k,j}\,(1 - F_{k,j}(t - t_k)) \\ &+ \sum_{j \textit{ (childs of } obs(k))} p_{k,j} F_{k,j}(t - t_k) p_{miss}(obs(j)) \\ &\left(\bar{p}_{obs(j)} + \sum_{l \textit{ (childs of } obs(j))} p_{j,l}\ (1 - F_{j,l}(t - t_k)) \right) \end{aligned}$$

The previous reasoning can be extended in a recursive way to include all of the observers down-stream from the current report's observer. The last term (between parenthesis) should be extended to include the probability of missed detection of the child observers of observer(j), and so on.

5 Multiple hypothesis tree management: the track forest

Tracks, both resolved and unresolved portions, are stored in a track forest. A forest is a structure that contains many trees, each with its own root node. Each node, other than a root node, is an instance of a report being hypothetically assigned to a track. Root nodes represent targets. A root node originally represents a new target starting with a single report. The root node is unresolved when it is created. When a root node is permanently declared as a new target, the root node becomes a resolved target node. Over time, as unresolved nodes become resolved, they are collapsed into the resolved root node of the target tree. The rest of the tree represents the unresolved portion of the target.

5.1 Track trees

Using Figure 3 as an illustration, we will explore the meaning of the track trees. The diagram is divided into four quadrants representing the trees after four reports are adde, one at a time. The portion of the tree created during the addition of the current report is shown in green. Report 1 (R1) is added as a root node in quadrant 1. The node has an identifier (a node ID), T1, or track 1. It is an unresolved root node of a target tree. In quadrant 2, assuming that report R2 gates with track T1, two nodes are added. Node T2 (track 2) represents a track that has R1 and R2 in it. Node T1 still represents the same thing as in the previous section, namely, a track with only one report, R1, in it. We also add the new unresolved target track root node T3, representing the possibility that report R2 was the beginning of a new target track. Continuing likewise in quadrant 3, we also assume that report R3 gates with all existing tracks. This will not be true in general, but we use this assumption for illustration of the track branching. R3 now appears in four new track nodes, including a new target in the node labelled track T6. Each track node where R3 appears represents a track that includes all ancestor nodes going back to the root. For example, the track T7 contains all three reports, R1, R2 and R3. Likewise, quadrant 4 shows the tree after adding report R4.

5.2 Global hypothesis list

A global hypothesis is a set of tracks satisfying the condition that all reports received by the tracker are assigned to exactly one track. Recall that each report generates a root unresolved target track node and a number of track nodes descended from existing track nodes. When a report is added to the track tree the global hypotheses are augmented and branched as follows:

1. *Branching step: For each track (known as the parent track) that the report gates with (including resolved tracks), and for all global hypotheses that the track was in, a new global hypothesis is created in which the child track replaces the parent track.*

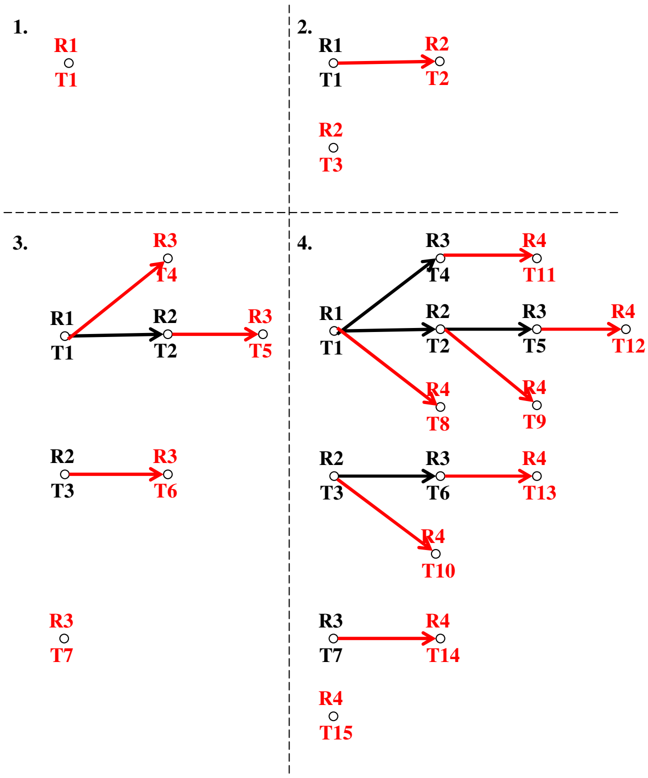

Figure 3: Track hypothesis tree growth as consecutive measurements are incorporated (the assumption here is that each report gates will all existing tracks, in a real situation this will not be true and the growth of the tree will be smaller than shown).

2. *Augmentation step: The new hypothetical root target track is added to all the original global hypotheses that were around before the branching step.*

We will now show the branching and augmentation steps for the same example shown in Figure (3). The tracker is initialized with a single empty global hypothesis. When the first report, R1 is added, there are no tracks in existence for it to gate with so the branching step is empty. The new unresolved target track, T1 is added to the single global hypothesis in the augmentation step. When the next report, R2, arrives and is added to the tree, the global hypothesis evolves into two global hypotheses. The first one is generated by adding a branch to the original global hypothesis (branching step). This hypothesis assumes that report R1 and R2 belong to the same target (track T2). The second hypothesis is generated by augmenting the original hypothesis with a new track T3 (augmentation step). This hypothesis assumes that reports R1 and R2 belong to two different targets (tracks T1 and T3). Since R3 gates with three existing tracks, each one in a single global hypothesis, three new global hypotheses are added in the branching step. The two pre-existing hypotheses are then augmented with the addition of track T6. Report R4 is then added using the same branching and augmentation algorithm, producing a total of fifteen global hypotheses. The above global hypothesis set

update process can be summarized as follows (where each global hypothesis is denoted by the tracks that belongs to it enclosed in {}):

1. *Report #1 added*
   - *Augmentation: {T1}*
2. *Report #2 added*
   - *Augmentation: {T1,T3}*
   - *Branching: {T2}*
3. *Report #3 added*
   - *Branching: {T4,T3}, {T1,T6}, {T5}*
   - *Augmentation: {T1,T3,T7}, {T2,T7}*
4. *Report #4 added*
   - *Branching: {T11,T3}, {T4,T10}, {T8,T6}, {T1,T13}, {T12} {T8,T3,T7}, {T1,T10,T7}, {T1,T3,T14}, {T9,T7}, {T2,T14}*
   - *Augmentation: {T4,T3, T15}, {T1,T6,T15}, {T5,T15} {T1,T3,T7,T15}, {T2,T7,T15}*

5.3 Global Hypothesis based pruning

It is well known, and also easy to see, that a Multiple Hypothesis Tracker like the one previously discussed needs some form of "hypothesis grow" control in order to be feasible to implement it in real time. In our case, we have implemented an adaptive pruning strategy that allows real time implementation while maximizing use of available cpu. The idea is to predict how long it will take to process a new measurement given the current track tree status. This is done after gating the new report with all feasible existing tracks, and deriving a relationship between the number of global hypothesis that need to be modified and the cpu time required for that. The parameters for this cpu time model are estimated on-line based on previous performance measurements. The predicted required time is then compared with the available time (computed based on the inter-report arrival time). If the predicted time is larger than the available time, then a fixed fraction of the less likely global hypotheses are discarded and all tracks that are not used by any of the remaining global hypotheses are pruned. This method has proven to be very robust. It has been tested in multiple cpu's with different processing speeds, and has always been able to maintain real time operation (over multiple days) with almost full utilization of available cpu (95% or higher). A different approach, based on track-tree pruning induced by resolving the oldest reports has shown lower performance than the global hypothesis based pruning.

6 Features and scores

Certain sensors are able to measure features in addition to the report's time. These features are modelled using a Gaussian state vector and Gaussian noise, and a standard Kalman filter is used to compute the measurement update likelihood. This likelihood is then combined with the kinematic likelihood to become the track likelihood.

But sometimes features are not well modelled by simple Gaussian models, and the "sensor" experts are more willing to compare two reports and give some kind of same/different measurement than to provide a Gaussian feature measurement and feature model. In our tracker formulation, we require that the answer to a comparison request between two reports be the likelihood that both reports originated from the same target ($p(\textit{report 1,report 2}|\textit{same})$) and the likelihood that both reports originated from different targets($p(\textit{report 1,report 2}|\textit{different})$). This new information regarding these two reports can now be used directly in the calculation of the track likelihood. For each global hypothesis, check whether both reports are in the same or different tracks (some global hypotheses will hypothesize that both reports are from the same target, and therefore both reports will be assigned to the same track, while other global hypotheses will assume that they are from different targets,and therefore the reports will be assigned to different tracks). The likelihood of every global where the reports are in the same track should be multiplied by the following likelihood ratio

$$L(\textit{report 1,report 2}) = \frac{p(\textit{report 1,report 2}|\textit{same})}{p(\textit{report 1,report 2}|\textit{different})}$$

This can be done in an asynchronous mode, where the tracker request scores between a new report and reports belonging to tracks that gated with the new reports. It then continues to process other reports, and when the score likelihoods become available, they are included into the likelihood of each global.

7 Conclusions

In this paper we have formulated a tracker that works with reports generated by a distributed network of sensors. The use of prior information (road network, transit times) as well as the use of negative information (probability of not-seen-yet) makes the formulation unique and quite different from other more common tracking problems like the radar or sonar tracking problems. Our approach to multiple hypothesis is also different from the most commonly used optimization-based techniques that estimate the single best global hypothesis. This tracker has been implemented in a real system with thirty sensors and up to five reports per second using a standard PC with acceptable results.

References

[1] *P. Arambel, M. Antonne, J. Douglas and R. Washburn,* ALPHATECH internal report TR-1135*, December 2002.*

[2] *C. Rago, C. Gershfield,* BAE Systems internal report TR-xxxx*, in preparation.*

[3] *A. Papoulis,* Probability, Random Variables and Stochastic Processes*, McGraw-Hill, Inc., 1991.*

[4] *Y. Bar-Shalom and X. Li,* Estimation and Tracking: Principles, Techniques and Software*, Artech House, 1993.*